REMARKS

Reconsideration is respectfully requested.

Rejection of Claims 1, 6, 11, 21 and 26 under 35 U.S.C. 112, second paragraph

The Office argues that claims 1, 6, 11, 21 and 26 are indefinite for failure to recite essential elements because there is no recitation of what happens when the two "if" clauses within each claim are not satisfied. Applicants respectfully traverse on the ground that these alternative cases are not essential. Both "if" clauses pertain to whether there is a prior version of a new data element (claims 1, 11 and 21) or a pointer forwarding entity (claims 6 and 26) that has been generated in a data element group. If there is a prior version, a version link is established between the new data element or pointer forwarding entity and its prior version (the action following the first "if" clause), and the prior version is freed following a grace period (the action following the second "if" clause). The absence of a prior version of the new data element requires no corresponding version linking or version freeing because there is no prior version to link to and no prior version to free. As such, it would be pointless to describe this trivial case in the rejected claims and applicants have properly chosen not to do so. Because claims 1, 6, 11, 21 and 26 are clear and definite on their face, and nothing essential to the claimed subject matter has been omitted, the rejection under 35 U.S.C. 112, second paragraph should be withdrawn.

Rejection of Claims 21-31 under 35 U.S.C. 101

Citing MPEP section 2106, the Office argues that claims 21-31 "lack a practical application of [sic, a] judicial exception (law of nature, abstract idea, naturally occurring article/phenomenon) since it fails to produce a useful result." The rejection further states that "the claimed subject matter fails to disclose a complete disclosure that contain [sic, contains]

some indication of the practical application for the claimed invention, i.e., why the applicant believes the claim is useful."

Applicants are somewhat baffled by this rejection because the first sentence of the rejection identities the very language that is said to be absent from the application, as follows:

> "These claims are directed to updating a shared data element group while preserving group integrity on behalf of one or more readers that are concurrently referencing group data elements without using locks or atomic instructions."

Applicants do not see what more is required than this statement (which is substantially recited in the preamble of independent claims 21, 26 and 31) to demonstrate that the subject matter of the rejected claims represents a practical application. Rather than recite a mere law of nature, abstract idea or naturally occurring article/phenomenon, the above-quoted statement confirms that the claims are directed to something that has demonstrative practical utility, namely the ability to update a shared data element group in a manner that preserves group integrity on behalf of readers without using expensive operations such as locks or atomic instructions. Because MPEP section 2106(II)(A) requires that "USPTO personnel should review the application to identify any asserted use," attention is further directed to the "Background of the Invention" section of the specification, and particularly page 5, lines 9-20, which describes in detail the prior art approach to maintaining data element group integrity on behalf of readers and the problems associated therewith, such as the need to lock the entire group while performing data updates or to copy the entire group if locking is not preformed. These problems are further discussed in the specification on page 5, line 21 to page 6, line 14 in the context of a cyclic search (such as within a finite state machine) as a specific example of a data element group requiring group integrity on behalf of readers. The claimed subject

matter solves these problems by eliminating the need for data element locking or group copying and instead provides an alternative solution that does not use these expensive operations and which allows readers to proceed unaffected through a data element group even while data elements are being updated. It will be readily seen from the foregoing sections of the application that the claimed subject matter is directed to solving a very real problem in a very real world environment, and that the claimed subject is not a mere law of nature, abstract idea or naturally occurring article/phenomenon divorced from a practical application.

MPEP section 2106(IV)(D) states that the Office has the burden of establishing a prima facie case of section 101 non-compliance that includes a duty to "identify and explain in the record the reasons why a claim is for an abstract idea with no practical application." MPEP section 2106(IV)(D) further states that "[i]f the record as a whole suggests that it is more likely than not that the claimed invention would be considered a practical application of an abstract idea, natural phenomenon, or law of nature, then USPTO personnel should not reject the claim."

Here, the Office has failed to identify which judicial exception applies to the rejected claims (i.e., law of nature, abstract idea, naturally occurring article/phenomenon) and the specific reasons why the exception applies. This is important because the Office may, for example, be interpreting claims 21, 26 and 31 as reciting a "mathematical algorithm" when in fact the meaning of this term has been questioned due to the fact that "any step-by-step process, be it electronic, chemical, or mechanical, involves an 'algorithm' in the broad sense of the term." *AT&T Corp. v. Excel Communications, Inc.*, 50 U.S.P.Q.2d 1447, 1450 (Fed. Cir. 1999). Thus, the required prima facie showing needed to support the section 101 rejection and to allow applicants an opportunity to fully respond has not been made.

Applicants cannot determine why the judicial exception argument is raised because nothing in the rejected claims suggests that such an exception applies. The fundamental utility inquiry under section 101 is whether the claimed subject matter does something "useful" instead of representing a mere idea. *State Street Bank & Trust Co. v. Signature Financial Group, Inc.*, 47 U.S.P.Q.2d 1596, 1601 (Fed. Cir. 1998) (". . . to be patentable, an algorithm must be applied in a 'useful' way"). Independent claims 21, 26 and 31 manifestly comply with this requirement by reciting a computer program product that programs a data processing platform to perform physical real world actions. For example, claim 21 recites programming means that program a data processing platform to perform operations such as generating a new group data element, assigning a generation number to the new data element, establishing a version link between the new data element and a prior version thereof, linking the new data element into the data element group, and freeing the prior version of the new data element. These actions represent manipulations of physical entities within the data processing platform (e.g., group element data structures) and are not simply unrealized abstractions. *See In re Lowry*, 32 U.S.P.Q. 2d 1031 (Fed. Cir. 1994) (memory resident data structures are physical entities). Indeed, the same actions set forth in rejected claim 21 are also recited in method claim 1, yet this claim has not been rejected under section 101 and has been determined to have utility. There is no principled reason to differentiate claim 21 from claim 1 when analyzing utility under section 101, for the analysis is the same regardless of the form – machine or process – in which a particular claim is drafted. *See AT&T, supra.* If it is the Office's position that computer program products are not patentable per se, applicants would remind the Office that this proposition was laid to rest over ten years ago when the BPAI withdrew its Federal Circuit appeal in *In re Beauregard*, 35 U.S.P.Q.2d 1383 (Fed. Cir. 1995) ("The Commissioner

now states that 'computer programs embodied in a tangible medium, such as floppy diskettes, are patentable subject matter under 35 U.S.C. §101 and must be examined under 35 U.S.C. §§ 102 and 103.'").

That the rejected claims 21-31 have practical utility is further demonstrated by applying the "physical transformation" and "useful, concrete and tangible result" tests of MPEP section 2106. The physical transformation test queries whether the claimed subject matter provides a transformation or reduction of an article to a different state or thing. In the rejected claims, the state of data element group is changed by creating a new element and linking it into the group in a prescribed fashion. Data transformations in a computer memory represent physical transformations that are useful. *See AT&T, supra. at 1452, State Street Bank, supra. at 1601.* The useful, concrete, tangible result test queries whether the claimed subject matter produces a result that is useful, concrete and tangible. Although it has been observed that these are not really separate elements requiring separate analysis, they are treated as such under MPEP 2106(IV)(C)2.(1)-(2). A result is said to be "useful" if it has specific, substantial and credible utility. As explained above, the rejected claims satisfy this element by improving the manner in which shared data element groups are updated while preserving group integrity on behalf of readers. Absent the claimed subject matter, readers would be slowed down due to the imposition of data element locking or memory requirements would be increased due to making a copy of the entire data element group. A result is said to be "concrete" if it is repeatable. This is clearly the case with the claimed subject matter. A result is said to be "tangible" if it is not abstract. As stated above, the rejected claims satisfy this requirement by modifying memory resident data structures in a physical (i.e., not abstract) manner. The processing efficiencies achieved by are also quite tangible.

Based on the foregoing, it would appear that claims 21-31 amply satisfy the utility requirements of section 101 and that the rejection based on this section should be withdrawn. As observed many years ago in *Diamond v. Diehr*, 209 U.S.P.Q. 1, 7-8 (U.S. 1981), the only things excluded from patent protection under the so-called judicial exception to section 101 are "laws of nature, physical phenomena and abstract ideas" claimed per se, and "[o]ur recent holdings in Gottshalk v. Benson, supra, and Parker v. Flook, supra, both of which are computer-related, stand for no more than these long established principles."

As a second ground for the section 101 rejection, the Office inconsistently argues that independent claims 21, 26 and 31 are directed to computer software per se, but then further argues that the addition of the medium recited in these claims does not satisfy section 101 requirements. The second argument refutes the first because the recitation of a medium precludes a finding that the claims are directed to computer software per se. It thus appears that the real ground for this portion of the section 101 rejection may be that claims 21, 26 and 31 are directed to nonfunctional descriptive material, which under MPEP section 2106.01[R-5], is not patentable even if it is associated with a medium. However, nonfunctional descriptive material under MPEP section 2106.01[R-5] refers to subject matter such as "music, literature, art, photographs, and mere arrangements or compilations of facts or data, without any functional interrelationship." Claims 21, 26 and 31 are not directed to this type of subject matter, and instead fall into the functional descriptive material category of MPEP section 2106.01[R-5], which includes data structures and programs. MPEP section 2106.01[R-5] states that this type of material satisfies section 101 if it is on a medium. Thus, claims 21, 26 and 31 should be deemed to overcome the section 101 for this reason as well.

As a third ground for the section 101 rejection, the Office argues that dependent claims 22-25 and 27-30 (1) recite "computing steps without producing any concrete and useful result and/or being limited to a practical application within the technological arts," (2) are "merely descriptive," and (3) "lack the necessary physical articles or objects to constitute a machine or a manufacture." The first argument was addressed above in the discussion demonstrating why the rejected claims are directed to a practical application rather than a law of nature, an abstract idea or a naturally occurring article/phenomenon. As discussed, the claimed subject matter produces a concrete and useful result by providing the ability to update a shared data element group in a manner that preserves group integrity on behalf of readers without using expensive operations such as locks or atomic instructions. This result is also tangible, thus satisfying the "useful, tangible, concrete" test for a "practical application" as set forth in MPEP section 2106(IV)(C)2.(2). The second argument relating to the claims being "merely descriptive" is not understood and seems to improperly borrow from trademark law. If the term "merely descriptive" is intended to mean that the claimed subject is not functional, applicants strenuously disagree because the sole purpose of the claimed subject matter is to produce a functional result by programming a data processing platform to take certain actions. Alternatively, if the Office is attempting to raise a "printed matter rejection," it is invited to consider *In re Lowry, supra.*, wherein this type of rejection was held to have no relevance to information stored in memory. The third argument relating to a lack of physical articles or objects to constitute a machine or a manufacture is wrong on its face because each of the claims recite a medium and are thus directed to a statutory article of manufacture. Therefore, claims 21-25 and 27-30 should be deemed to overcome the section 101 for these reasons as well.

Obviousness Rejection of Claims 1-30 under 35 U.S.C. 103

Claims 1-30 are rejected for obviousness as being unpatentable over the so-called "Applicant Admitted Prior Art" ("APA") set forth in the specification at pages 1-6 and in Figs. 1A-3, taken in view of the McKenney article entitled "Read-Copy Update." In dealing with independent claims 1, 6, 11, 16, 21 and 26, the Office puts forth two rejections, namely, a "first rejection" based on the claims with their "if" clauses removed and a "second rejection" based on the claims with the "if" clauses intact. As demonstrated above in connection with the section 112 rejection, the "if" clauses do not render the claims indefinite. Even they did, it would be improper to simply eliminate these clauses from the claims for purposes of analysis under section 103. Such an analysis (i.e., the first rejection) is based on hypothetical subject matter divorced from any real claim set forth in the application. The remainder of this discussion will thus respond to the "second rejection."

Independent Claims 1, 11 and 21

Claims 1, 11 and 21 respectively recite a method, a system and a computer program product for updating a "shared data element group" so as to allow updates thereof "while preserving group integrity on behalf of one or more readers that are concurrently referencing group data elements without using locks or atomic instructions." McKenney at page 1, lines 12-16 is argued as disclosing this limitation. However, this passage merely discusses the fact that conventional read-copy update preserves the old state of an individual data element being updated. This does not guarantee the preservation of "group integrity," which, as discussed in the APA on page 5, lines 9-20, requires that group updates be implemented atomically.

Claims 1, 11 and 21 each further recite "assigning a generation number to said new data element that allows a reader of said data element group to determine whether said new

data format is a correct version for said reader." McKenney at page 15, column 1, section 6.1, paragraph 1 is argued as disclosing this limitation. However, this passage is describing per-CPU generation sequence numbers that are used to track read-copy update grace periods. The generation sequence numbers are not assigned to new data elements of a data element group so that readers can determine the correct data element version. Rather, the generation sequence numbers are assigned to CPUs and are used to identify which of a pair of per-CPU counters is to be manipulated by readers that preempt or block while inside their critical sections. One per-CPU counter is for a current generation of readers and the other per-CPU counter is for a previous generation of readers. When the previous generation counter is reset to zero, the per-CPU generation sequence number is advanced to create a new generation and the roles of the per-CPU counters are reversed. The current generation counter becomes the previous generation counter and the previous generation counters becomes the current counters. After the generation sequence number on a given CPU has been advanced two times, it is deemed that all operations in existence prior to the advancement have terminated on that CPU, i.e., that a grace period has expired. Thus, McKenney fails to disclose or suggest the recited claim limitation of assigning a generation number to a new data element that allows a reader of the data element group to determine whether the new data format is a correct version for the reader.

Claims 1, 11 and 21 each further recite 'if a version of said new data element exists, establishing a version link between said new data element and said prior version." The APA at page 3, lines 5-10 is argued as disclosing this limitation. However, this passage is discussing Fig. 1C and the fact that the stale data element B retains a linked list pointer to data element C for the benefit of a reader r1 that is currently referencing data element B after a new

data element B' has been created. This is not the claimed "version link between said new data element and said prior version." To disclose the claim limitation, there would have to be a linking entity such as a pointer from the new data element B' to the old data element B, but this is not shown or described insofar as that is part of the claimed invention and not the prior art being discussed in Fig. 1C.

Claims 1, 11 and 21 each further recite "updating a global generation number associated with said data element group." McKenney at page 15, column 1, section 6.1 is argued as disclosing this limitation, along with Fig. 29 and page 16, #3. The section 6.1 passage and the significance of the per-CPU generation sequence numbers were discussed above. The generation sequence numbers are used for grace period tracking and are assigned to CPUs in order to identify which of two per-CPU counters is to be incremented by readers. The generation sequence numbers are not associated with a data element group as claimed. Fig. 29 of McKenney does not alter this conclusion. This figure illustrates data structures used by a read-copy update subsystem to accumulate callbacks and to track grace periods that determine when the callbacks may be processed. Page 16, #3 of McKenney is assumed to be directed to the third element of "Item 5." Here, the discussion is about global current-generation counters, maximum generation counters, and per-CPU generation counters, all of which are used to track grace periods for callback processing purposes and none of which are associated with a data element group.

Based on the foregoing, it is respectfully asserted that independent claims 1, 11 and 21 are patentably distinguishable over the APA and McKenney art.

Independent Claims 6, 16 and 26

Claims 6, 16 and 26 respectively recite a method, a system and a computer program product for updating a "shared data element group" so as to allow updates thereof "while preserving group integrity on behalf of one or more readers that are concurrently referencing group data elements without using locks or atomic instructions." McKenney at page 3, line 23 – page 4, line 4 is cited as disclosing this limitation. However, the cited passage merely discusses a conventional "Toy" implementation of conventional read-copy update.

Claims 6, 16 and 26 each further recite "assigning a generation number to said pointer forwarding entity that allows a reader of said data element group to determine whether said pointer forwarding entity is a correct version for said reader." McKenney at page 15, column 1, section 6.1, paragraph 1 is argued as disclosing this limitation. However, as discussed above, this passage is describing per-CPU generation sequence numbers that are used to track read-copy update grace periods. The generation sequence numbers are not assigned to pointer forwarding entities of a data element group so that readers can determine the correct pointer forwarding entity version.

Claims 6, 16 and 26 each further recite "if there is a prior version of said pointer forwarding entity, establishing a version link between said pointer forwarding entity and said prior version." The APA at page 3, lines 5-10 is argued as disclosing this limitation. However, as discussed above, this passage is discussing Fig. 1C and the fact that the stale data element B retains a linked list pointer to data element C for the benefit of a reader r1 that is currently referencing data element B after a new data element B' has been created. This is not the claimed version link between a pointer forwarding entity and a prior version.

Claims 6, 16 and 26 each further recite "updating a global generation number associated with said data element group." McKenney at page 15, column 1, section 6.1 is argued as disclosing this limitation, along with Fig. 29 and page 16, #3. See the discussion above for applicants' rebuttal to this aspect of the rejection, which was also raised in connection with claim 1, 11 and 21.

Independent Claims 7, 17 and 27

Claims 7, 17 and 27 each recite "assigning a current global generation number to said search." McKenney at page 16, column 2, #3 is argued as disclosing this limitation. However, none of the generation counters mentioned in this passage represent a current global generation number that is assigned to a search. Rather, the counters are used to track grace periods for callback processing.

Claims 7, 17 and 27 each further recite "when referencing a data element in said data element group, determining whether said referenced data element is a correct version by comparing a generation number assigned to said referenced data element with said search generation number." McKenney at page 15, column 1, paragraph 4 is argued as disclosing this limitation. However, as discussed above, this passage is describing per-CPU generation sequence numbers that are used to track read-copy update grace periods. The generation sequence numbers are not assigned to data elements and are not referenced by readers to determine whether a referenced data element is a correct version by comparing the generation number assigned to the data element with a search generation number.

Claims 7, 17 and 27 each further recite "searching for a correct version of said referenced data element as necessary." McKenney at page 17, column 1, paragraph 4 is argued as disclosing this limitation. However, this passage refers to a CPU advancing its

callback lists to facilitate callback processing when a global generation number is advanced. This does not occur as part of searching a data element group.

Dependent Claims

The remaining dependent claims are believed to be allowable based on the independent claims from which they depend. Additional grounds for allowance of the dependent claims include the following:

Claims 3, 13 and 23: The references do not disclose or suggest the claimed subject matter wherein the subject matter of claims 1, 11 and 21 is respectively used "to delete a group data element and said new data element is generated by copying said data element to be deleted and setting a deletion flag in said new data element." McKenney at page 7, column 1, paragraph 4 mentions data element deletion and copying in the same paragraph, but does not disclose copying as part of a deletion operation. McKenney at page 7, column 1, paragraph 3 mentions flagging stale data so that a reader will know the data is stale, but there is no disclosure of a delete flag used for identifying a data element as deleted.

Claims 5, 15 and 25: The references do not disclose or suggest the claimed subject matter wherein the subject matter of claims 1, 11 and 25 respectively further includes "generating a pointer-forwarding entity that points to said new data element, said pointer forwarding entity maintaining said version link on behalf of said new data element and further being used to link said new data element into said data element group." The APA at page 3, lines 1-5 does not disclose this feature. Rather, this passage is directed Fig. 1C and describes how a new data element is created by copying an old data element and linking the new data element into a linked list. There is no disclosure of the claimed pointer forwarding entity.

Claims 8, 18 and 28: The references do not disclose or suggest the claimed subject that respectively depends from claims 7, 17 and 27, and "wherein, if said data element generation number is equal to said search generation number, said referenced data element is accepted for reading as a correct version." McKenney at page 17, column 1, paragraphs 3 and 4 is directed to a CPU's manipulation of generation counters to advance a grace period and move callback lists for callback processing. There is no disclosure of accepting a data element for reading if it is a correct version of a referenced data element.

Claims 9, 19 and 29: The references do not disclose or suggest the claimed subject matter that respectively depends from claims 7, 17 and 27, and "wherein, if said data element generation number is less than said search generation number, a search is made for a later version of said referenced data element, and wherein said referenced data element is used if a later version is not found." McKenney at page 5, column 1, paragraph 3 is directed to conventional read-copy update searching without locks. There is no disclosure of searching for a later version of a referenced data element if a data element generation number is less than a search generation number.

Claims 10, 20 and 30: The references do not disclose or suggest the claimed subject matter that respectively depends from claims 7, 17 and 27, and "wherein, if said data element generation number is greater than said search generation number, a search is made for a prior version of said referenced data element, and wherein said referenced data element is deemed to be a new insertion if there is no prior version. Mckenney at page 18, column 1, paragraph 2 is directed to a reader-writer lock that uses per-CPU locks that are cache-aligned for rapid acquisition by readers. There is no disclosure of searching for a prior version of a data

element if it has a data element generation number that is greater than a search generation number.

Anticipation Rejection of Claim 31 under 35 U.S.C. 102(a)

The APA is said to disclose all of the elements of claim 31. However, there are several claim limitations that are entirely absent in the APA. For example, the APA does not discloses "A computer program product for managing a shared data element group so as to allow updates thereof while preserving group integrity on behalf of one or more readers that are concurrently referencing group data elements without using locks or atomic instructions." The APA discloses only a conventional read-copy update that allows lock free read operations while individual data elements are being updated but does not 'þreserve group integrity on behalf of one or more readers that are concurrently referencing group data elements without using locks or atomic instructions." As set forth at the end of the APA, this is the problem to be solved by the claimed invention.

The APA also fails to disclose 'þerforming a first-phase update operation that preserves a consistent pre-update view of said data element group on behalf of pre-update readers and a consistent post-update view of the data element group on behalf of post-update readers." Again, the APA does not disclose maintaining group consistency, except as a statement of the problem to be solved. The passage at page 2, lines 1-8 is referring to pre-update and post-update views of individual data elements, not an entire data element group.

The APA also fails to disclose 'þroviding means by which readers can locate all data elements of said data element group that belong to each of said pre-update and post-update views as readers search said data element group." The conventional read-copy update technique disclosed in the APA does not allow readers to locate all data elements that belong to pre-update

and post-update views. The view that is preserved by conventional read-copy update is the data element currently being referenced by the reader. The passage at page 2, lines 3-6 merely confirms this aspect of convention read-copy update.

The APA also fails to disclose 'performing one or more read operations following said first-phase update operation in which one or more readers search said data element group with each reader referencing only data elements belonging to one of said pre-update and post-update views." The passage at page 2, lines 17-19 merely refers to the fact that conventional read-copy update allows multiple readers to concurrently traverse a linked list without locks.

Based on the foregoing, it is respectfully requested that the anticipation rejection of claim 31 be withdrawn.

In view of the foregoing, Applicants respectfully request that all rejections be withdrawn and that Notices of Allowability and Allowance be issued.

Respectfully submitted,

/Walter W. Duft/

Walter W. Duft
Attorney for Applicant
Registration No. 31,948

Law Offices of Walter W. Duft
8616 Main Street, Suite 2
Williamsville, New York 14221
Telephone: (716) 633-1930
Facsimile: (716) 633-1932